arXiv:1612.05623v1 [cs.DS] 16 Dec 2016

Approximating Approximate Distance Oracles

Michael Dinitz
Department of Computer Science
Johns Hopkins University

Zeyu Zhang
Department of Computer Science
Johns Hopkins University

Abstract

Given a finite metric space (V, d), an approximate distance oracle is a data structure which, when queried on two points $u, v \in V$, returns an approximation to the the actual distance between u and v which is within some bounded stretch factor of the true distance. There has been significant work on the tradeoff between the important parameters of approximate distance oracles (and in particular between the size, stretch, and query time), but in this paper we take a different point of view, that of per-instance optimization. If we are given an particular input metric space and stretch bound, can we find the smallest possible approximate distance oracle for that particular input? Since this question is not even well-defined, we restrict our attention to well-known classes of approximate distance oracles, and study whether we can optimize over those classes.

In particular, we give an $O(\log n)$-approximation to the problem of finding the smallest stretch 3 Thorup-Zwick distance oracle, as well as the problem of finding the smallest Pătraşcu-Roditty distance oracle. We also prove a matching $\Omega(\log n)$ lower bound for both problems, and an $\Omega(n^{\frac{1}{k}-\frac{1}{2^{k-1}}})$ integrality gap for the more general stretch $(2k-1)$ Thorup-Zwick distance oracle. We also consider the problem of approximating the best TZ or PR approximate distance oracle *with outliers*, and show that more advanced techniques (SDP relaxations in particular) allow us to optimize even in the presence of outliers.

1 Introduction

Given a finite metric space (V, d), an approximate distance oracle is a data structure which can approximately answer distance queries. It is usually a combination of a preprocessing algorithm to compute a data structure, and a query algorithm which returns a distance $d'(u, v)$ whenever queried on a pair of vertices $u, v \in V$. An approximate distance oracle is said the have *stretch* t if $d(u, v) \leq d'(u, v) \leq t \cdot d(u, v)$. Note that there is a trivial stretch 1 distance oracle that uses $\Theta(n^2)$ space: we could just store the entire metric space. So the goal is to reduce the space, i.e., to build a small data structure that also has small stretch and small query time.

The seminal work on approximate distance oracles is due to Thorup and Zwick [21]. They showed that for every integer $k \geq 1$, every finite metric space has an approximate distance oracle with stretch $(2k - 1)$ and query time $O(k)$ which uses only $O(kn^{1+\frac{1}{k}})$ space. A significant fraction of more recent results have built off of the ideas developed in [21], and much of this follow-up work has stored the exact same (or very similar) data structure, just with improved query algorithms or slightly different information in the storage (see, e.g., [18, 23, 7, 8]). Most notably, Pătraşcu and Roditty [17] gave a different distance oracle (still using some of the basic ideas from [21]) that has multiplicative stretch of 2 and additive stretch of 1, with size $O(n^{\frac{5}{3}})$. This broke through the stretch 3 barrier from [21]. Later this result was improved to more general multiplicative/additive stretches [1].

In this paper we ask a natural but very different type of question about approximate distance oracles: can we find (or approximate) the *best* approximate distance oracle? If we are given an input metric space and a stretch bound, is it possible to find the *smallest* approximate distance oracle for that particular input? This is an unusual question in two ways. First, most data structures are by design forced to store all of the input data; the question is *how* to store it and what *extra* information should be stored. This is the case in other settings where instance-optimality of data structures has been considered, e.g., static or dynamic optimality of splay trees. Second, it is not clear whether this question is even well-defined: lower bounds on data structures are commonly arrived at through information or communication complexity (see, e.g., [16]) but when we ask for the optimal data structure on one particular instance this approach becomes meaningless.

However, approximate distance oracles are different in ways which allow us to make meaningful progress towards these optimization questions. First, since we are allowed to return only approximate distances (up to some stretch factor), we are allowed to store only part of the input (and indeed this is the entire point of such an oracle). The second problem is a bit more tricky: given an input, how can we optimize over "the space of all approximate distance oracles"? What does this mean, and what does this space look like?

To get around this issue, we make an observation: many modern distance oracles (and in particular Thorup-Zwick, Pătraşcu-Roditty, and almost all of their variants) have a similar structure. The preprocessing algorithm chooses a subset of the original distances to store which has some particular structure, and the query algorithm can return a valid distance estimate efficiently as long as the stored distances satisfy the required structure. Thus we can optimize for *these particular* distance oracles by choosing the *best possible* set of distances to remember subject to the required structure. By characterizing this structure for different types of distance oracles, we can optimize over those types.

For example, the stretch-3 Thorup-Zwick distance oracle uses a subtle but simple method to choose the set of distances to store. It randomly samples a subset of approximately $\sqrt{n}$ vertices, without using any information about the original metric space, and then creates a data structure which is related (in a well-defined, important way) to these vertices. The correctness of the query algorithm does not depend on the choice of the vertices. Thus instead of simply choosing the

subset of vertices uniformly at random, we can instead try to optimize the set of chosen vertices with respect to the actual input metric space.

In this paper, we give matching $\Theta(\log n)$ upper and lower bounds for optimizing stretch-3 Thorup-Zwick distance oracles, and matching $\Theta(\log n)$ upper and lower bounds for optimizing the Pătraşcu-Roditty distance oracle. These upper bounds both use a similar LP relaxation, but by giving an $\Omega(n^{\frac{1}{k}-\frac{1}{2^{k-1}}})$ integrality gap for optimizing stretch-$(2k-1)$ Thorup-Zwick distance oracles, we show that this relaxation is not enough to give nontrivial approximations when extended to larger stretch values.

As an extension, we also study the problem of optimizing distance oracles *with outliers*: if we are allowed to not answer queries for some of the vertices (of our choosing), can we have much smaller storage space? We give an $(O(\log n), 1+\varepsilon)$-bicriteria approximation to both stretch-3 Thorup-Zwick and Pătraşcu-Roditty distance oracles with outliers. We also give a true approximation to stretch-3 Thorup-Zwick distance oracle with outliers when the number of outliers is small.

Relationship to Spanners. It is worth noting that this paper is motivated by a similar line of research on *graph spanners* (subgraphs which approximately preserve distances). Spanners and distance oracles tend to be related (although there is no known formal connection between them), and the traditional questions asked of spanners (what is the tradeoff between the stretch and the size?) are similar to the traditional questions asked of distance oracles. Recently, there has been significant progress in looking at spanners from an optimization point of view: given an input graph and an allowed stretch bound, can we find the sparsest possible spanner meeting that stretch bound? In the last few years, upper and lower bounds have been developed for these problems in the basic case, the directed case, with a degree objective, with fault-tolerance, etc. See, e.g., [11, 4, 12, 10, 9].

It is natural to ask these kinds of optimization questions for distance oracles as well, but the definitions become much more difficult. For spanners, the space we are optimizing over (all subgraphs) is very clear and well-defined. But for distance oracles, as discussed, it is much harder to define the space of all data structures. Thus in this paper we optimize over restricted classes, where this space is more well-defined. We view our definitions of these restricted optimization questions as one of the major contributions of this work.

2 Definitions and Preliminaries

We begin with some basic definitions, including formal definitions of the problems that we will be working on.

Definition 2.1. An approximate distance oracle with (m, a)-stretch, size s, preprocessing time g, and query time h is a pair of algorithms, *preprocess* and *query*, with the following properties.

- *preprocess* is a randomized preprocessing algorithm $preprocess(V, d, m, a, r)$ which takes as input a metric space (V, d), stretch bound (m, a), and random string r and outputs a data structure S where the expected output size is at most $\mathbb{E}_r[|S|] \leq s(|V|, m, a)$ and the expected preprocessing time is at most $g(|V|, m, a)$.
- *query* takes as input a data structure $S = preprocess(V, d, m, a, r)$ (the output of the preprocess algorithm) with two vertices $u, v \in V$, and outputs a value $d'(u, v) \in \mathbb{R}$ such that $d(u, v) \leq d'(u, v) \leq m \cdot d(u, v) + a$. The running time of *query* is at most $h(|V|, m, a)$.

We will frequently refer to these just as "distance oracles" rather than "approximate distance oracles" when the stretch bound is clear from context.

The query algorithm guarantees here are deterministic: the randomness only affects the size of the data structure. Note that one could easily define distance oracles so that either the correctness (with respect to the stretch bound) or the query running time (or both) hold only in expectation or with high probability, but as discussed in Section 1, essentially all existing distance oracles (and in particular the Thorup-Zwick distance oracle) have deterministic guarantees on the queries.

This naturally leads us to the following question: If we fix a particular distance oracle and metric space, can we find the *best* possible data structure? Here we will focus on the output size, not the preprocessing time (as long as the preprocessing time is polynomial). In other words, since the query algorithm work on *any* of the possible data structures which the preprocessing algorithm might output, can we actually find the smallest such data structure? This gives the following natural optimization problem.

Definition 2.2. Given an approximate distance oracle $\mathcal{A} = (preprocess, query)$, the $\mathcal{A}$-optimization problem takes as input a metric space (V, d) and a stretch bound (m, a), and the goal is to find a string r which minimizes $|preprocess(V, d, m, a, r)|$.

In this paper we will focus on two distance oracles (Thorup-Zwick [21] and Pătraşcu-Roditty [17]), so we now introduce these oracles.

2.1 Thorup-Zwick Distance Oracle

For every integer $k \geq 1$, Thorup and Zwick [21] provided an approximate distance oracle with $(2k-1, 0)$-stretch, size $O(n^{1+\frac{1}{k}})$, preprocessing time $O(kn^{2+\frac{1}{k}})$, and query time $O(k)$. We call this distance oracle TZ_k.

Their preprocessing algorithm first constructs a chain of subsets $\varnothing = A_k \subseteq A_{k-1} \subseteq \ldots \subseteq A_0 = V$ by repeated sampling. Each set A_i, where $i \in [k-1]$, is obtained by including each element of A_{i-1} independently with probability $n^{-\frac{1}{k}}$.

Let $R_{iu} = \{v \in A_{i-1} \mid d(u,v) < \min_{w \in A_i} d(u,w)\}$ for all $u \in V$ and $i \in [k]$ (where by convention $\min_{w \in \varnothing} d(u,w) = \infty$ for all $u \in V$ to handle the $i = k$ case). The output data structure is obtained by storing (in a 2-level hash table) the distance from each node u to each node in $\bigcup_{i=1}^{k} R_{iu}$.

The data structure also stores a little more information. Each vertex u remembers $k-1$ pivots: $\arg\min_{w \in A_i} d(u,w)$ for all $i \in [k-1]$, and the distance from u to these pivots. However, this is a fixed space cost, and also negligible, so when analyzing the size of the oracle we will ignore the cost of storing the pivots

Clearly the output data structure is determined once $A_1, \ldots, A_{k-1}$ are fixed. The size of the data structure is:

$$cost(A_1, \ldots, A_{k-1}, V, d) = \sum_{u \in V} \sum_{i=1}^{k} |R_{iu}| = \sum_{u \in V} \sum_{i=1}^{k} \left| \{v \in A_{i-1} \mid d(u,v) < \min_{w \in A_i} d(u,w)\} \right|.$$

We will refer to $\sum_{u \in V} |R_{iu}|$ as the cost in level i.

Let us look back on the definition of approximate distance oracle. The random string r is only used to generate A_i's, and the query algorithm will return a correct distance estimate no matter what the sets A_i are, but the size is determined by the sets. Therefore, the TZ_k-optimization problem is to find the subsets $\varnothing = A_k \subseteq A_{k-1} \subseteq \ldots \subseteq A_0 = V$ in order to minimize the total cost.

2.2 Pătraşcu-Roditty Distance Oracle

Pătraşcu and Roditty [17] provided an approximate distance oracle with $(2,1)$-stretch, size $O(n^{\frac{5}{3}})$, preprocessing time $O(n^2)$, and query time $O(1)$. We call this distance oracle PR. Note that PR works only for metric spaces with integer distances.

Their preprocessing algorithm first construct a set $A \subseteq V$ via a complicated correlated sampling (informally, they sample a large set and a small set, and then define A to be everything in the large set and everything contained in a ball around the small set delimited by the large set). The data structure consists of a 2-level hash table for the distance from each node in A to each node in V, as well as a 2-level hash table storing the distance between each pair $\{u, v\} \subseteq V$ such that $d(u,v) < \min_{w \in A} d(u,w) + \min_{w \in A} d(v,w) - 1$

As with Thorup-Zwick, the output data structure is completely determined once A is fixed. Let $R = \{\{u,v\} \subseteq V \mid d(u,v) < \min_{w \in A} d(u,w) + \min_{w \in A} d(v,w) - 1\}$. Then the size of the data structure is

$$cost(A, V, d) = n \cdot |A| + |R| = n \cdot |A| + \left| \left\{ \{u,v\} \subseteq V \mid d(u,v) < \min_{w \in A} d(u,w) + \min_{w \in A} d(v,w) - 1 \right\} \right|.$$

As before, the random string r is only used to generate the set A, and any $A \subseteq V$ gives a data structure on which the query algorithm works. Therefore, the PR-optimization problem is to find the subset $A \subseteq V$ in order to minimize the total cost.

2.3 Distance Oracles With Outliers

In some cases, a small set of outlier vertices may make the size of the data structure blow up. Yet in some applications it is acceptable to ignore these outliers. This was the motivation behind a line of work on distance oracles with slack ([5], [6]), in which the data structure could ignore the stretch bound on a small fraction of the distances.

In this paper, we consider the case that we can refuse to answer distance queries for some outlier vertices. In other words, we can essentially remove an outlier set F out of V when computing the distance oracle. This gives us the problem of optimizing distance oracle with outliers, in which we not only need to find the random string to determine the output data structure, we also need to find the set of outliers to minimize the final cost. More formally, we have the following type of problem.

Definition 2.3. Given an approximate distance oracle $\mathcal{A} = (preprocess, query)$, the $\mathcal{A}$-optimization problem with outliers takes as input a metric space (V,d), a stretch bound (m,a), and a bound on the number of outliers $f \in \mathbb{N}$. The goal is to find a string r as well as a set $F \subseteq V$ where $|F| \leq f$, in order to minimize $|preprocess(V \backslash F, d, m, a, r)|$.

We will provide both true approximation results and (α, β)-bicriteria results, in which we slightly violate the bound on the number of outliers. Formally, an (α, β)-approximation algorithm for the $\mathcal{A}$-optimization problem with outliers is on algorithm which on any input $((V,d),(m,a),f)$ returns a solution with cost at most $\alpha \cdot OPT$ that has at most $\beta \cdot f$ outliers (where OPT is the minimum cost of any solution with at most f outliers).

2.4 Our Results and Techniques

With these definitions in hand, we can now formally state our results.

In Section 3 we discuss the problem of optimizing the 3-stretch Thorup-Zwick distance oracle, i.e., the TZ_2-optimization problem. It is straightforward to obtain an $O(\log n)$-approximation by reducing to the non-metric facility location problem.

Theorem 2.4. *There is an $O(\log n)$-approximation algorithm for the TZ_2-optimization problem.*

To prove a matching lower bound, we use a reduction from Label Cover to the TZ_2-optimization problem. We use a proof which is similar to the proof of the hardness of Set Cover in [22] (based on [14]). However, we cannot use a reduction directly from Set Cover since we will need some extra properties of the starting instances, and thus are forced to start from Label Cover. We introduce a new notion of (m, l, δ)-set families and show that these can still be plugged into existing hardness results to get the extra structural properties that we need. This lets us prove the following theorem:

Theorem 2.5. *Unless* $\mathbf{NP} \subseteq \mathbf{DTIME}(n^{O(\log \log n)})$*, the TZ_2-optimization problem does not admit a polynomial-time $o(\log n)$-approximation.*

For larger stretch values, a natural approach is to realize that a simple LP relaxation suffices to give Theorem 2.4 in the stretch 3 case, and try to extend this basic LP to larger stretches. In Section 4, we show that this does not work for the more general TZ_k-optimization problem: the integrality gap jumps up to become a polynomial. The instance is very simple: it is just the metric space formed by shortest paths on the n-cycle. It turns out to be straightforward to calculate the optimal fractional LP cost, but proving that the optimal integral solution is large is surprisingly involved.

Theorem 2.6. *The basic LP relaxation for the TZ_k-optimization problem has an $\Omega(n^{\frac{1}{k}-\frac{1}{2^k-1}})$ integrality gap when $k > 2$.*

In Section 5 we discuss the problem of optimizing the Pătraşcu-Roditty distance oracle. The basic LP and a simple rounding algorithm gives us an $O(\log n)$-approximation algorithm.

Theorem 2.7. *There is an $O(\log n)$-approximation algorithm for PR-optimization problem.*

A reduction from set cover problem also gives us a matching lower bound.

Theorem 2.8. *Unless* $\mathbf{P} = \mathbf{NP}$*, the PR-optimization problem does not admit a polynomial-time $o(\log n)$-approximation.*

In Section 6 we move to the outliers setting. For both TZ_2- and PR-optimization problems, a semidefinite programming relaxation and a simple rounding algorithm gives us an $(O(\frac{\log n}{\varepsilon}), 1+\varepsilon)$-approximation algorithm. Using an SDP relaxation seems to be necessary – the corresponding LP relaxation requires violating the number of outliers by a factor of 2 rather than a factor of $1+\varepsilon$. We can also get a true approximation on TZ_2-optimization problem with outliers if the number of outliers is low. These results form the following theorems.

Theorem 2.9. *There is an $(O(\frac{\log n}{\varepsilon}), 1+\varepsilon)$-approximation algorithm for the TZ_2-optimization problem with outliers.*

Theorem 2.10. *There is an $O(\log n)$-approximation algorithm for TZ_2-optimization problem with outliers if the number of outliers is at most $\sqrt{n}$.*

Theorem 2.11. *There is an $(O(\frac{\log n}{\varepsilon}), 1+\varepsilon)$-approximation algorithm for the PR-optimization problem with outliers.*

3 TZ_2-Optimization Problem

We first give an $O(\log n)$-approximation for TZ_2-optimization (Theorem 2.4), and follow this with a matching lower bound.

3.1 Upper Bound

We will prove our upper bound by a reduction to the non-metric facility location problem.

Definition 3.1. In the *non-metric facility location* problem we are given a set F of facilities, a set D of clients, an opening cost function $f : F \to \mathbb{R}^+$, and a connection cost function $c : D \times F \to \mathbb{R}^+$. The goal is to find the set $S \subseteq F$ which minimizes $\sum_{i \in S} f(i) + \sum_{i \in D} \min_{j \in S} c(i, j)$ (i.e. the sum of the opening and connection costs).

Non-metric facility location is a classic problem, and much is known about it, including the following upper bound due to Hochbaum.

Theorem 3.2 ([15])**.** *There is a polynomial time algorithm which gives an $O(\log n)$-approximation to the non-metric facility location problem.*

Hochbaum's algorithm is a greedy algorithm, but it is also straightforward to design an algorithm with similar bounds using an LP relaxation. Since it is not necessary we do not present the relaxation here, but generalizations of the relaxation will prove important in the more general TZ_k setting (see Section 4).

We now show that the TZ_2-optimization problem is essentially a special case of non-metric facility location problem. First, simple arithmetic manipulation of the cost function of the TZ_2-optimization problem gives the following:

$$\begin{aligned}
cost(A_1, V, d) &= \sum_{u \in V} |R_{1u}| + \sum_{u \in V} |R_{2u}| \\
&= \sum_{u \in V} \left| \{v \in V \mid d(u, v) < \min_{w \in A_1} d(u, w)\} \right| + \sum_{u \in V} |\{v \in A_1 \mid d(u, v) < \infty| \\
&= \sum_{u \in V} \left| \{v \in V \mid d(u, v) < \min_{w \in A_1} d(u, w)\} \right| + n|A_1| \\
&= \sum_{w \in A_1} n + \sum_{u \in V} \min_{w \in A_1} |\{v \in V \mid d(u, v) < d(u, w)\}| \,.
\end{aligned}$$

Given an instance (V, d) of the TZ_2-optimization problem, we create an instance of non-metric facility location by setting $F = D = V$, opening costs $f(v) = n$ for all $v \in V$, and connection costs $c(u, w) = |\{v \in V \mid d(u, v) < d(u, w)\}|$ for all $u, w \in V$. Then the cost function of the TZ_2-optimization problem is exactly the same as the cost function of non-metric facility location problem. Therefore TZ_2 is a special case of non-metric facility location, which together with Theorem 3.2 implies Theorem 2.4.

3.2 Lower Bound

Proving an $\Omega(\log n)$ hardness of approximation (Theorem 2.5) turns out to be surprisingly difficult. Details appear in Appendix A; here we provide an informal overview. Technically we reduce directly to TZ_2-optimization from a version of the Label Cover problem that corresponds to applying parallel repetition [19] to 3SAT-5, which is a standard starting point for hardness reductions. Informally, though, we are "really" reducing from Set Cover: given an instance of Set Cover, we show how to create an instance of TZ_2-optimization where the cost of the optimal solution is the same (up to a constant and a polynomial scaling factor). But in order for our reduction to work, we actually need more than just an arbitrary Set Cover instance: we need a version of Set Cover in which it is hard even to cover *most* of the elements, not just all of them.

So we have to also give a new reduction from Label Cover to Set Cover, showing that even this version of Set Cover is hard. It turns out that Feige's reduction [14], reinterpreted by Vazirani [22], essentially already gives us what we need. We just need to analyze it a bit more carefully. In particular, a key component of this reduction is what Vazirani called (m, l)-*set systems*, which can be thought of as nearly-unbiased sample spaces. We generalize this notion to (m, l, δ)-set systems, given in the following definition.

Definition 3.3. A set B (the universe) and a collection of subsets $C_1, \ldots, C_m$ of B form an *(m, l, δ)-set system* if any collection of l sets in $\{C_1, \ldots, C_m, \overline{C_1}, \ldots, \overline{C_m}\}$ whose union contains at least $(1 - \delta)|B|$ elements must include both C_i and $\overline{C_i}$ for some i.

An (m, l)-set system is just a $(m, l, 0)$-set system. While not all (m, l)-set systems are $(m, l\delta)$-set systems for larger δ, the construction of (m, l)-set systems in [22] actually does generalize directly to larger values of δ. With this tool in hand, we follow through the rest of the reduction and it gives us the type of Set Cover instances which we need. Technically our reduction skips this step by going directly from Label Cover to TZ_2-optimization, but generating these kinds of Set Cover instances is intuitively what the first part of the reduction is doing.

4 TZ_k-Optimization Problem

We now move to the more general TZ_k-optimization problem. While we are not able to give nontrivial upper bounds for this problem, we can at least show that the basic LP relaxation (as discussed in Section 3.1) does not give polylogarithmic bounds in this more general setting.

4.1 The LP

Let $B_u(v) = \{w \in V \mid d(u, w) \leq d(u, v)\}$. For every $v \in V$ and $i \in [k]$, let $x_v^{(i)}$ be a variable which is supposed to be an indicator for whether $v \in A_i$. Similarly, for all $u, v \in V$ and $i \in [k]$, let $y_{uv}^{(i)}$ be a variable which is supposed to be an indicator for whether $v \in R_{iu}$. (Recall that $R_{iu} = \{v \in A_{i-1} \mid d(u, v) < \min_{w \in A_i} d(u, w)\}$) We can easily write an LP relaxation for this problem:

$$
\begin{array}{rlll}
(LP_{TZ_k}): \min & \sum_{i=1}^{k} \sum_{u,v \in V} y_{uv}^{(i)} & & \\
s.t. & 0 = x_v^{(k)} \leq x_v^{(k-1)} \leq \ldots \leq x_v^{(1)} \leq x_v^{(0)} = 1 & & \forall v \in V \\
& y_{uv}^{(i)} \geq x_v^{(i-1)} - \sum_{w \in B_u(v)} x_w^{(i)} & & \forall u, v \in V, i \in [k] \\
& y_{uv}^{(i)} \geq 0 & & \forall u, v \in V, i \in [k]
\end{array}
$$

It can easily be shown that this is a valid relaxation (the proof can be found in Appendix B). When restricted to the special case of $k = 2$, it is not hard to see that this LP is essentially a special case of the basic LP relaxation for non-metric facility location, which can be used to prove the $O(\log n)$ bound of Theorem 2.4. But for larger values of k the behavior is different, and does not result in a polylogarithmic integrality gap.

4.2 Integrality Gap

The integrality gap instance is quite simple: the metric (V, d) induced by shortest-path distances in a cycle. Slightly more formally, we let $V = [n]$, and use the cycle distance $d(u, v) = \min\{|u - v|, n + \min\{u, v\} - \max\{u, v\}\}$.

Details can be found in Appendix B. It turns out to be relatively easy to find a fractional solution to LP_{TZ_k} with cost $O(n^{1+\frac{1}{2^k-1}})$ on this instance. The tricky part is lower bounding the optimal solution, i.e., showing that the optimal integral solution has cost at least $\Omega(n^{1+\frac{1}{k}})$. Combining these two results gives us an $\Omega(n^{\frac{1}{k}-\frac{1}{2^k-1}})$ integrality gap, proving Theorem 2.6.

5 PR-Optimization Problem

We now move from Thorup-Zwick distance oracles to Pătraşcu-Roditty distance oracles. We show that from an optimization perspective, they are similar to TZ_2 in that we can find matching bounds: an $O(\log n)$-approximation, and $\Omega(\log n)$-hardness.

5.1 Upper Bound

In this section we prove Theorem 2.7 by using an LP and randomized rounding to give an $O(\log n)$-approximation to the PR-optimization problem.

Let $B_u(v) = \{w \in V \mid d(u,w) \leq d(u,v)\}$, and $B(u,r) = \{w \in V \mid d(u,w) \leq r\}$. We can see $B_u(v) = B(u, d(u,v))$. Now, let x_v be a variable which is supposed to be an indicator for whether $v \in A$, and let y_{uv} be a variable which is supposed to be an indicator for whether $\{u,v\} \in R$. (Recall that $R = \{\{u,v\} \subseteq V \mid d(u,v) < \min_{w\in A} d(u,w) + \min_{w\in A} d(v,w) - 1\}$). We can write the following LP relaxation:

$$\begin{array}{lll} (LP_{PR}): \min & \sum_{v\in V} n \cdot x_v + \sum_{\{u,v\}\subseteq V} y_{uv} & \\ s.t. & y_{uv} \geq 1 - \sum_{w\in B_u(r)\cup B_v(d(u,v)-r)} x_w & \forall u,v \in V, \forall r \in [0, d(u,v)] \\ & x_v \in [0,1] & \forall v \in V \\ & y_{uv} \geq 0 & \forall u,v \in V \end{array}$$

At first blush it may not be obvious that the first type of constraint in this LP really captures the characterization of paris in R. But it is actually not that hard to see that this is a valid relaxation (a formal proof can be found in Appendix C). Note that while the number of constraints appears to be exponential (recall that we assume integer weights, but not necessarily unit weights, and hence $d(u,v)$ is not necessarily polynomial in the input size), it is in fact possible to solve this LP in polynomial time. We can do this by noting that for each $u,v \in V$, only at most n different value of r actually yield *different* constraints, so we can simply write the constraints for those values.

Our algorithm is relatively straightforward. We first solve LP_{PR} and get an optimal fractional solution (x^*_v, y^*_{uv}). We then use independent randomized rounding, adding each $v \in V$ to A independently with probability $\min\{4\ln n \cdot x^*_v, 1\}$.

Lemma 5.1. *If $y^*_{uv} \leq \frac{1}{2}$, then the probability that $\{u,v\} \in R$ is at most $\frac{1}{n}$.*

Proof. If $y^*_{uv} \leq \frac{1}{2}$, then the first constraint implies that $\sum_{w\in B_u(r)\cup B_v(d(u,v)-r)} x^*_w \geq \frac{1}{2}$ for all $r \in [0, d(u,v)]$. Therefore, the probability that $A \cap (B_u(r) \cup B_v(d(u,v)-r)) = \varnothing$ for a specific $r \in [0, d(u,v)]$ is at most

$$\prod_{w\in B_u(r)\cup B_v(d(u,v)-r)} (1 - \min\{4\ln n \cdot x^*_w, 1\}) \leq e^{-\sum_{w\in B_u(r)\cup B_v(d(u,v)-r)} 4\ln n \cdot x^*_w} \leq \frac{1}{n^2}$$

A union bound over all the different values of r we used in our LP implies that the probability that there exists an $r \in [0, d(u,v)]$ where $A \cap (B_u(r) \cup B_v(d(u,v)-r)) = \varnothing$ is at most $\frac{1}{n^2} \cdot n = \frac{1}{n}$.

We claim that the existence of such an r is implied by $\{u,v\} \in R$, and hence the probability that $\{u,v\} \in R$ is at most $\frac{1}{n}$. To see this, suppose that $\{u,v\} \in R$, i.e. suppose that $d(u,v) < \min_{w \in A} d(u,w) + \min_{w \in A} d(v,w) - 1$. Then if we set $r = \min_{w \in A} d(u,w) - 1$, this implies that $\min_{w \in A} d(v,w) > d(u,v) - r$. But then this would imply that no element of A is in $B_u(r) \cup B_v(d(u,v) - r)$. □

Let $OPT_{LP_{PR}}$ denote the optimal cost of LP_{PR}. Then the above lemma implies that the expected cost of the rounding algorithm is at most

$$\begin{aligned}\mathbf{E}[n|A| + |R|] &\leq \sum_{v \in V} n \cdot x_v^* \cdot 4 \ln n + 2 \cdot \sum_{u,v \in V} y_{uv}^* + n^2 \cdot \frac{1}{n} \leq O(\log n) \cdot OPT_{LP_{PR}} + n \\ &\leq O(\log n) \cdot OPT\end{aligned}$$

(where we use the fact that $OPT \geq \Omega(n)$). This completes the proof of Theorem 2.7.

5.2 $\Omega(\log n)$-hardness

We now show a matching hardness bound for the PR-optimization problem by reducing from the Set Cover problem.

Consider a set cover instance $(\mathcal{U}, \mathcal{S})$ where $|\mathcal{U}| + |\mathcal{S}| = n$. For each $e \in \mathcal{U}$, we create a group of vertices G_e where $|G_e| = 3n$. For each $S \in \mathcal{S}$, we also create a group of vertices G_S where $|G_S| = 3n$.

Now we construct the following metric space: $V = (\bigcup_{e \in \mathcal{U}} G_e) \cup (\bigcup_{S \in \mathcal{S}} G_S)$ and

$$d(u,v) = \begin{cases} 1, & \text{if } u \in G_e, v \in G_e \\ 1, & \text{if } u \in G_S, v \in G_S \\ 1, & \text{if } u \in G_e, v \in G_S, e \in S \\ 2, & \text{otherwise.} \end{cases}$$

In Appendix C we show that if there is a solution $\mathcal{S}^*$ to the set cover instance $(\mathcal{U}, \mathcal{S})$ where $|\mathcal{S}^*| = t$, then there is a set A where $cost(A, V, d) \leq t|V|$. We also show that if there is a set $A \subseteq V$ where $cost(A, V, d) \leq t|V|$, then there exists a solution $\mathcal{S}^*$ to the set cover instance $(\mathcal{U}, \mathcal{S})$ where $|\mathcal{S}^*| = t$. These two claims, together with an appropriate hardness theorem for Set Cover [13], imply Theorem 2.8.

6 Distance Oracles With Outliers

We now move to the more difficult outliers setting, where we can also optimize over a set of vertices to ignore. Recall that for an approximate distance oracle $\mathcal{A}$, our goal is now to find a set of vertices F (the outliers) where $|F| \leq f$ as well as a string r in order to minimize $|preprocess(V \setminus F, d, m, a, r)|$. In other words, we are going to try to solve the same problems as before, but where we can choose a set F to remove. We begin with TZ_2, and then move to PR.

6.1 TZ_2-Optimization Problem With Outliers

For this problem, it is easy to see that the cost function becomes:

$$cost(A, F, V, d) = (n-f)|A| + \sum_{u \in V \setminus F} |R_{1u}| = (n-f)|A| + \sum_{u \in V \setminus F} \left| \{v \in V \setminus F \mid d(u,v) < \min_{w \in A} d(u,w)\} \right|$$

A natural approach is to use an LP which is similar to LP_{TZ_k} to solve this problem (but for TZ_2), suitably adapted to handle outliers. Let x_v be a variable which is supposed to be an indicator for whether $v \in A$, let y_{uv} be a variable which is supposed to be an indicator for whether $v \in R_{1u}$, and let z_v be a variable which is supposed to be an indicator for whether $v \in F$. Then we can write the following natural LP relaxation:

$$\begin{array}{rll}
(LP_{TZ_2O}): \min & \sum_{v\in V}(n-f)\cdot x_v + \sum_{u,v\in V} y_{uv} & \\
s.t. & y_{uv} \geq 1 - z_u - z_v - \sum_{w\in B_u(v)} x_w & \forall u,v \in V \\
& \sum_{v\in V} z_v \leq f & \\
& x_v \in [0,1] & \forall v \in V \\
& y_{uv} \geq 0 & \forall u,v \in V \\
& z_v \in [0,1] & \forall v \in V
\end{array}$$

Unfortunately, this LP can not give an (α, β)-approximation with $\beta = 2 - \epsilon$. To see this, consider the case that $f = \frac{n}{2}$. Then the optimal solution to LP_{TZ_2O} is 0, by setting all z_v to $\frac{1}{2}$, all x_v to 0, and all $y_{uv} = 0$. Thus any integral solution, to be competitive with this fractional solution, must treat *all* nodes as outliers, requiring β to be at least 2.

Fortunately we can give a stronger semidefinite programming relaxation, allowing for a better approximation. As in LP_{TZ_2O}, let $\vec{x}_v$ be a variable which is supposed to be an indicator for whether $v \in A$, let $\vec{y}_{uv}$ be a variable which is supposed to be an indicator for whether $v \in R_{1u}$, and let $\vec{z}_v$ be a variable which is supposed to be an indicator for whether $v \in F$. We can then write this SDP:

$$\begin{array}{rll}
(SDP_{TZ_2O}): \min & \sum_{v\in V}(n-f)\cdot \|\vec{x}_v\|^2 + \sum_{u,v\in V} \|\vec{y}_{uv}\|^2 & \\
s.t. & \|\vec{y}_{uv}\|^2 \geq 1 - \vec{z}_u \cdot \vec{z}_v - \sum_{w\in B_u(v)} \|\vec{x}_w\|^2 & \forall u,v \in V \\
& \sum_{v\in V} \|\vec{z}_v\|^2 \leq f & \\
& \|\vec{x}_v\|^2 \leq 1 & \forall v \in V \\
& \|\vec{y}_{uv}\|^2 \leq 1 & \forall u,v \in V \\
& \|\vec{z}_v\|^2 \leq 1 & \forall v \in V
\end{array}$$

Our approximation algorithm first solves SDP_{TZ_2O} to get an optimal solution $(\vec{x}^*_v, \vec{y}^*_{uv}, \vec{z}^*_v)$. We then use independent randomized rounding to construct A, adding each $v \in V$ to A independently with probability $\min\{\frac{3\ln n}{\varepsilon} \cdot \|\vec{x}^*_v\|^2, 1\}$ where ε is a small constant. Finally, we use threshold rounding to construct F by adding each $v \in V$ to F if $\|\vec{z}^*_v\|^2 \geq \frac{1}{1+\varepsilon}$.

We want to show that this is an $(O(\log n), 1+\varepsilon)$-approximation. It is easy to see that $|F| \leq (1+\varepsilon)f$ because $\sum_{v\in V} \|\vec{z}^*_v\|^2 \leq f$. In order to prove Theorem 2.9 it only remains to prove that the expected cost is at most $O(\log n) \cdot OPT$. This proof can be found in Appendix D.

When $f \leq \sqrt{n}$ we can actually give a true $O(\log n)$-approximation (Theorem 2.10). The algorithm is almost the same; we just need to change the threshold rounding for outliers to instead pick the f vertices with largest $\|\vec{z}_v\|^2$ value. Details appear in Appendix D.

6.2 PR-Optimization Problem With Outliers

For this problem, the cost function becomes:

$$\begin{aligned}
cost(A,F,V,d) =& (n-f)\cdot|A| + |R| \\
=& (n-f)\cdot|A| + \left|\{\{u,v\} \subseteq V\backslash F \mid d(u,v) < \min_{w\in A} d(u,w) + \min_{w\in A} d(v,w) - 1\}\right|.
\end{aligned}$$

We will again use an SDP relaxation. Let $\vec{x}_v$ be a variable which is supposed to be an indicator for whether $v \in A$, let $\vec{y}_{uv}$ be a variable which is supposed to be an indicator for whether $\{u,v\} \in R$,

and let $\vec{z}_v$ be a variable which is supposed to be an indicator for whether $v \in F$. We have the following relaxation which is similar to both LP_{PR} and SDP_{TZ_2O}:

$$\begin{array}{rll}
(SDP_{PR}): \min & \sum_{v\in V}(n-f)\cdot\|\vec{x}_v\|^2 + \sum_{\{u,v\}\subseteq V}\|\vec{y}_{uv}\|^2 & \\
s.t. & \|\vec{y}_{uv}\|^2 \geq 1 - \vec{z}_u\cdot\vec{z}_v - \sum_{w\in B_u(r)\cup B_v(d(u,v)-r)}\|\vec{x}_w\|^2 & \forall u,v\in V, r\in[0,d(u,v)] \\
& \sum_{v\in V}\|\vec{z}_v\|^2 \leq f & \\
& \|\vec{x}_v\|^2 \leq 1 & \forall v\in V \\
& \|\vec{y}_{uv}\|^2 \leq 1 & \forall u,v\in V \\
& \|\vec{z}_v\|^2 \leq 1 & \forall v\in V
\end{array}$$

Note that this SDP_{PR} is solvable in polynimial time for the same reason that LP_{PR} is solvable: for each pair of (u,v), we can find n different values of r that give all of the distinct constraints.

The rounding algorithm is basically the same as the TZ_2-optimization problem with outliers. We first solve the SDP_{PR} and get an optimal solution $(\vec{x}^*_v, \vec{y}^*_{uv}, \vec{z}^*_v)$. We then use independent randomized rounding to get A, adding each $v \in V$ to A independently with probability $\min\{\frac{6\ln n}{\varepsilon}\cdot\|\vec{x}^*_v\|^2, 1\}$ where ε is a small constant. Then we use threshold rounding to get F, adding each $v \in V$ to F if $\|\vec{z}^*_v\|^2 \geq \frac{1}{1+\varepsilon}$.

This is an $(O(\log n), 1+\varepsilon)$-approximation. It is easy to see that $|F| \leq (1+\varepsilon)f$ because $\sum_{v\in V}\|\vec{z}^*_v\|^2 \leq f$. The proof that the expected cost is at most $O(\log n)\cdot OPT$ is in Appendix D, which completes the proof of Theorem 2.11.

7 Conclusion and Future Work

In this paper we initiate the study of *approximating* approximate distance oracles. This is a different take on the question of optimizing data structures, where we attempt to find the best data structure for a particular input, rather than for a class of inputs. In order to make this tractable (or even well-defined), we restrict our attention to known classes of distance oracles, and show that it is sometimes possible to find the best of these restricted oracles. We also extended our approaches to optimize in the presence of outliers.

For future work, the major question is clearly whether we can approximately optimize higher level (i.e., higher stretch) Thorup-Zwick distance oracles. Although we show an integrality gap for the basic LP, it is quite conceivable that a stronger LP or SDP could be used to give a logarithmic approximation ratio. Beyond this, there are other distance oracles which could be optimized – we chose Thorup-Zwick and Pătraşcu-Roditty since they are well-known and in some ways canonical, but it would be interesting to extend these ideas to other oracles. At a higher level, we believe that the definitions and ideas we have introduced here could lead to many interesting questions about optimizing data structures for given inputs: can we find near-optimal distance labels? Or compact routing schemes? Or connectivity oracles? Or fault-tolerant oracles? Essentially any data structure question in which there is a choice of *which* data to store, rather than how to store it, can be put into our optimization framework. Exploring this space is an exciting future direction.

References

[1] Ittai Abraham and Cyril Gavoille. On approximate distance labels and routing schemes with affine stretch. In *International Symposium on Distributed Computing*, pages 404–415. Springer, 2011.

[2] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.

[3] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 113–131. ACM, 1988.

[4] Piotr Berman, Arnab Bhattacharyya, Konstantin Makarychev, Sofya Raskhodnikova, and Grigory Yaroslavtsev. Approximation algorithms for spanner problems and directed steiner forest. *Information and Computation*, 222:93–107, 2013.

[5] T-H Hubert Chan, Kedar Dhamdhere, Anupam Gupta, Jon Kleinberg, and Aleksandrs Slivkins. Metric embeddings with relaxed guarantees. *SIAM Journal on Computing*, 38(6):2303–2329, 2009.

[6] T-H Hubert Chan, Michael Dinitz, and Anupam Gupta. Spanners with slack. In *European Symposium on Algorithms*, pages 196–207. Springer, 2006.

[7] Shiri Chechik. Approximate distance oracles with constant query time. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 654–663. ACM, 2014.

[8] Shiri Chechik. Approximate distance oracles with improved bounds. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 1–10. ACM, 2015.

[9] Eden Chlamtác and Michael Dinitz. Lowest degree k-spanner: Approximation and hardness. In *APPROX-RANDOM*, volume 28 of *LIPIcs*, pages 80–95. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.

[10] Eden Chlamtac, Michael Dinitz, and Robert Krauthgamer. Everywhere-sparse spanners via dense subgraphs. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 758–767. IEEE Computer Society, 2012.

[11] Michael Dinitz and Robert Krauthgamer. Directed spanners via flow-based linear programs. In *Proceedings of the forty-third annual ACM Symposium on Theory of computing*, pages 323–332. ACM, 2011.

[12] Michael Dinitz and Zeyu Zhang. Approximating low-stretch spanners. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 821–840. SIAM, 2016.

[13] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 624–633. ACM, 2014.

[14] Uriel Feige. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.

[15] Dorit S Hochbaum. Heuristics for the fixed cost median problem. *Mathematical programming*, 22(1):148–162, 1982.

[16] Guy Joseph Jacobson. Succinct static data structures. 1988.

[17] Mihai Patrascu and Liam Roditty. Distance oracles beyond the thorup-zwick bound. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 815–823. IEEE, 2010.

[18] Mihai Patrascu, Liam Roditty, and Mikkel Thorup. A new infinity of distance oracles for sparse graphs. In *Foundations of Computer Science (focs), 2012 Ieee 53rd Annual Symposium on*, pages 738–747. IEEE, 2012.

[19] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998.

[20] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability pcp characterization of np. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 475–484. ACM, 1997.

[21] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.

[22] Vijay V Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.

[23] Christian Wulff-Nilsen. Approximate distance oracles with improved query time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 539–549. Society for Industrial and Applied Mathematics, 2013.

A Proofs in section 3 (lower bound of TZ_2-optimization problem)

A.1 Label Cover Problem

For the lower bound, we start from a hard Label Cover instance, and following the steps in proving the hardness of approximating Set Cover problem. Since the definition of the Label Cover problem is somewhat complex, we break it into parts: first defining an instance, a labelling, and then defining the problem. Note that we are using a specific setting where the parameters in the graph are strongly related, so it is slightly different from the definition of classic/general Label Cover problem.

Definition A.1. A label cover instance consists of $(G = (V_1, V_2, E), \Sigma, \Pi)$ where

- G is a bipartite graph between vertex sets V_1 and V_2 and an edge set E. Let $V' = V_1 \cup V_2$
- G is left and right regular. Denote by Δ_1 and Δ_2 the degrees of vertices in V_1 and V_2 respectively.
- For each edge e, there is a constraint Π_e which is a bijection function from Σ to itself. The set of all constraints in G are $\Pi = \{\Pi_e : \Sigma \to \Sigma \mid e \in E\}$

Definition A.2. A *labelling* of the graph, is a mapping $\sigma : V' \to \Sigma$ which assigns a label for each vertex of G. A labelling σ is said to satisfy an edge $e = (v_1, v_2)$ if and only if $\Pi_e(\sigma(v_1)) = \sigma(v_2)$.

The following definition is the problem which will be the starting point of our reduction.

Definition A.3. In the LabelCover$_{n,r,\varepsilon}$ problem, we are given an instance (G, Σ, Π) of Label Cover where $|V_1| = (5n)^r, |V_2| = (5n)^r, |\Sigma| = 7^r, \Delta_1 = 15^r, \Delta_2 = 15^r$, and one of the following is true:

- There exists a labelling σ such that it satisfies all the edges e in G (in which case we say that the input is a YES instance), or
- For any labelling σ of the vertices, no more than $\varepsilon^r|E|$ edges are satisfied by σ (in which case we say that the input is a NO instance).

The goal is to determine whether the input is a YES or a NO instance.

Label Cover forms the starting point of many hardness of approximation results. Its hardness is a now-classical application of the PCP theorem [3] and Raz's parallel repetition lemma [19], which give the following theorem.

Theorem A.4 ([19])**.** *There exists a constant $\varepsilon \geq 0$ such that LabelCover$_{n,r,\varepsilon}$ is not in* **P** *unless* **NP** $\subseteq$ **DTIME**$(n^{O(r)})$.

For example, there exists a constant $\varepsilon \geq 0$ such that LabelCover$_{n,3\log\log n,\varepsilon}$ is not in **P** unless **NP** $\subseteq$ **DTIME**$(n^{O(\log\log n)})$. Starting from here, we will fix $r = 3\log\log n$, and ε be the constant which makes LabelCover$_{n,3\log\log n,\varepsilon}$ hard.

A.2 (m, l, δ)-Set System

We also need a (m, l, δ)-set system (see Definition 3.3) to do the reduction. We can construct a (m, l, δ)-set system by using a (l, γ)-independent collection of length m strings.

Definition A.5. Let B be a collection (may contains repetitions) of binary strings of length m. B is (l,γ)-*independent* if the following inequality holds for every $i_1, i_2, \ldots, i_l$ and $a \in \{0,1\}^l$:

$$\left| \Pr_{x \in B}[x_{i_1} = a_1 \wedge \ldots \wedge x_{i_l} = a_l] - 2^{-l} \right| \leq \gamma.$$

A corollary of lemma 1 and construction 3 in [2] provides a explict construction of (l,γ)-independent collection.

Corollary A.6. *For any $l \leq m$, there is an explict construction of a $(l, (1-2^{-l})\cdot 2^{-l-1})$-independent collection of length m strings with $|B| = 4^{l+1}m^2$ in $|B|^{O(1)}$ time.*

With the corollary in hand, we can construct a (m,l,δ)-set system with the parameters we want.

Lemma A.7. *For any $l \leq m$, there is an explicit construction of a $(m, l, 2^{-l-1})$-set system with $|B| = 4^{l+1}m^2$ in $|B|^{O(1)}$ time.*

Proof. Let B be the collection of length m strings in Corollary A.6. Define $C_i = \{x \in B \mid x_i = 1\}$ for all $i \in [m]$, we will show that $(B; C_1, \ldots, C_m)$ is a $(m, l, 2^{-l-1})$-set system.

Assume that there exist $D_{i_1}, D_{i_2}, \ldots, D_{i_l}$ such that $\left|\bigcup_{j=1}^{l} D_{i_j}\right| \geq (1-2^{-l-1})|B|$, where each D_{i_j} is either C_{i_j} or $\overline{C_{i_j}}$ (note that this implies that there are no j and k such that $D_{i_j} = \overline{D_{i_k}}$). Define

$$a_j = \begin{cases} 0, & \text{if } D_{i_j} = C_{i_j} \\ 1, & \text{if } D_{i_j} = \overline{C_{i_j}} \end{cases}.$$

Let $S = \{x \mid x \in B, x_{i_1} = a_1, x_{i_2} = a_2, \ldots, x_{i_l} = a_l\}$. Because B is a $(l, (1-2^{-l})\cdot 2^{-l-1})$-independent collection, we have

$$|S| = |\{x \mid x \in B, x_{i_1} = a_1, x_{i_2} = a_2, \ldots, x_{i_l} = a_l\}| > (2^{-l} - (1-2^{-l})\cdot 2^{-l-1})|B| > 2^{-l-1}|B|.$$

On the other hand, note by construction, for all $x \in S$ and $j \in [l]$, we have $x \notin D_{i_j}$, which implies that $\left|\bigcup_{j=1}^{l} D_{i_j}\right| \leq |B| - |S| < (1-2^{-l-1})|B|$: a contradiction. □

A.2.1 Reduction

We now show how to use the set systems from the previous section to give a reduction from Label Cover to TZ_2-optimization problem.

Let $(G = (V_1, V_2, E), \Sigma, \Pi)$ be a LabelCover$_{n,r,\varepsilon}$ instance with $r = 3\log\log n$, and let $(B; C_1, \ldots, C_m)$ be a $(m, l, 2^{-l-1})$-set system with $m = |\Sigma| = 7^r, l = r\log n$.

We first create a universe $\mathcal{U} = E \times B$, and a set of sets $\mathcal{S} = \{S_{v,x} \mid v \in V', x \in [m]\}$ (recall that $V' = V_1 \cup V_2$). Here

$$S_{v,x} = \bigcup_{e: v \in e, e \in E} \{e\} \times C_{\Pi_e(x)}, \text{ if } v \in V_1,$$

$$S_{v,x} = \bigcup_{e: v \in e, e \in E} \{e\} \times \overline{C_x}, \text{ if } v \in V_2.$$

We know that $|E| = (15n)^r$, $|B| = 4^{r\log n+1}\cdot 7^{2r} = n^{\Theta(1)\cdot r}$ and $|\mathcal{S}| = m\cdot|V'| = 7^r\cdot 2\cdot(5n)^r$, so $|\mathcal{U}| \gg |\mathcal{S}|$. Without lose of generality and for simplicity of our proof, we can assume $|\mathcal{U}|$ is dividable

by $|\mathcal{S}|$, so that we can replicate $\mathcal{S}$ for $\frac{|\mathcal{U}|}{|\mathcal{S}|}$ times, and get a set of sets $\mathcal{S}' = \mathcal{S}^{(1)} \cup \ldots \cup \mathcal{S}^{\left(\frac{|\mathcal{U}|}{|\mathcal{S}|}\right)}$ which has the same size as $\mathcal{U}$.

It is also easy to see that each $u = ((v_1, v_2), b) \in \mathcal{U}$ appears in exactly m sets in $\mathcal{S}$ because for each $x \in [m]$, either $u \in S_{v_1,x}$ or $u \in S_{v_2,\Pi_{(v_1,v_2)}(x)}$. Therefore each $u \in \mathcal{U}$ appears in $\frac{m|\mathcal{U}|}{|\mathcal{S}|} = \frac{|\mathcal{U}|}{|V'|}$ sets in $\mathcal{S}'$.

The metric space is defined as $V = \mathcal{U} \cup \mathcal{S}'$ and the distance is defined as following:

$$d(u,v) = \begin{cases} 1.2, & \text{if } u \in \mathcal{S}', v \in \mathcal{S}' \\ 1.4, & \text{if } u \in v \text{ or } v \in u \\ 1.6, & \text{if } u \in \mathcal{U}, v \in \mathcal{U} \\ 1.8, & \text{otherwise} \end{cases}$$

This metric space (V, d) will form the instance of TZ_2-optimization which we analyze. It is easy to see that the reduction is polynomial because $|V|$ is polynomial of $|E|$.

A.2.2 Analysis

Lemma A.8. *If (G, Σ, Π) is a YES instance in the LabelCover$_{n,r,\varepsilon}$ problem. Then the reduction (V, d) to the TZ_2-optimization problem has a solution with cost $\leq (|V'| + 1) \cdot |V|$.*

Proof. Let $\sigma : V' \to [m]$ denote a labelling of G which satisfies all the edges in E. Let $A_1 = \{S_{v,\sigma(v)} \mid v \in V'\}$. We claim that A_1 is a solution with cost at most $(|V'| + 1) \cdot |V|$.

Note that in Section 3.1 we showed that the level 2 cost $\sum_{u \in V} |R_{2u}| = |V| \cdot |A_1| = |V'| \cdot |V|$, the only thing left is to show that the level 1 cost is $\sum_{u \in V} |R_{1u}| \leq |V|$. We will prove this by showing $R_{1u} \subseteq \{u\}$ for all $u \in V$.

For any $u \in \mathcal{S}'$, we have that $d(u, v) \geq 1.2$ for all $v \in V$ because the definition of d, and $\min_{w \in A_1} d(u, w) = 1.2$ because $A_1 \cap \mathcal{S}' \neq \varnothing$. Thus $R_{1u} = \{v \in V \mid d(u, v) < \min_{w \in A_1} d(u, w)\} \subseteq \{u\}$.

For any $u = ((v_1, v_2), b) \in \mathcal{U}$, we have that $d(u, v) \geq 1.4$ for all $v \in V$ because the definition of d. We also know that either $u \in S_{v_1,\sigma(v_1)}$ or $u \in S_{v_1,\Pi_{(v_1,v_2)}(\sigma(v_1))}$ by the definition of $\mathcal{S}$, and $\Pi_{(v_1,v_2)}(\sigma(v_1)) = \sigma(v_2)$ because edge (v_1, v_2) is satisfied by labelling σ. Therefore $u \in S_{v_1,\sigma(v_1)} \cup S_{v_2,\sigma(v_2)}$. From the fact that both $S_{v_1,\sigma(v_1)}$ and $S_{v_2,\sigma(v_2)}$ are in A_1, we have $\min_{w \in A_1} d(u, w) = 1.4$. Thus $R_{1u} = \{v \in V \mid d(u, v) < \min_{w \in A_1} d(u, w)\} \subseteq \{u\}$.

Therefore $R_{1u} \subseteq \{u\}$ for all $u \in V$, so that A_1 is a solution with cost at most $\leq (|V'|+1) \cdot |V|$. □

Lemma A.9. *If (G, Σ, Π) is a No instance in the LabelCover$_{n,r,\varepsilon}$ problem. Then the reduction (V, d) to the TZ_2-optimization problem has no solution with cost $< \frac{l}{8}|V'| \cdot |V|$.*

Proof. We prove the lemma by showing that if the optimal solution of the reduction (V, d) to the TZ_2-optimization problem has cost $< \frac{l}{8}|V'| \cdot |V|$, then there exists a labelling σ such that it satisfies more than $\varepsilon^r |E|$ edges.

Assume the optimal solution $A_1 \subseteq V$ has $cost(A_1, V, d) < \frac{l}{8}|V'| \cdot |V|$, then $|A_1| < \frac{l}{8}|V'|$ because the level 2 cost is $\sum_{u \in V} |R_{2u}| = |V||A_1|$.

Let $L_v = \{x \in [m] \mid \exists j \in \left[\frac{|\mathcal{U}|}{|\mathcal{S}|}\right] \text{ s.t. } S^{(j)}_{v,x} \in A_1 \cap \mathcal{S}'\}$ for all $v \in V$, then $\sum_v L_v \leq |A_1 \cap \mathcal{S}'| \leq |A_1| < \frac{l}{8}|V'|$. Therefore at least $\frac{3}{4}|V'|$ vertices has $|L_v| < \frac{l}{2}$, because otherwise $\sum_v L_v \geq (1 - \frac{3}{4}) \cdot |V'| \cdot \frac{l}{2} \geq \frac{l}{8}|V'|$.

Let $E_1 = \{e = (v_1, v_2) \in E \mid |L_{v_1}| < \frac{l}{2}, |L_{v_2}| < \frac{l}{2}\}$. Then $|E_1| \geq \frac{|E|}{2}$ because $|V_1| = |V_2| = \frac{|V'|}{2}$ and G is regular.

On the other hand, we define a $u \in \mathcal{U}$ is "uncovered" if $\{v \in A_1 \cap \mathcal{S}' \mid u \in v\} = \varnothing$. Then for any uncovered $u \in \mathcal{U}$, we know that $\min_{w \in A_1} d(u, w) = 1.6$. Thus

$$\begin{aligned} R_{1u} =& \{v \in V \mid d(u, v) < \min_{w \in A_1} d(u, w)\} \\ \geq& \{v \in \mathcal{S}' \mid d(u, v) < 1.6\} \\ =& \{v \in \mathcal{S}' \mid u \in v\} = \frac{|\mathcal{U}|}{|V'|}. \end{aligned}$$

Therefore $|\{u \in \mathcal{U} \mid u \text{ is uncovered}\}| < \frac{\frac{l}{8}|V'| \cdot |V|}{\frac{|\mathcal{U}|}{|V'|}} < \frac{l}{4}|V'|^2$.

Let $E_2 = \{e \in E \mid |\{u = (e, b) \in \mathcal{U} \mid u \text{ is uncovered}\}| < \frac{l|V'|^2}{|E|}\}$. Then $|E_2| \geq \frac{3}{4}|E|$ because otherwise $|\{u \in \mathcal{U} \mid u \text{ is uncovered}\}| \geq (|E| - \frac{3}{4}|E|) \cdot \frac{l|V'|^2}{|E|}\} \geq \frac{l}{4}|V'|^2$.

Let $E' = E_1 \cap E_2$, we know that $|E'| \geq \frac{|E|}{4}$.

Now, we will show that if we uniformly random sample labels from L_v for each $v \in V'$, the expected number of the edges satisfied in E' is at least $\frac{|E|}{l^2}$.

For each edge $e = (v_1, v_2) \in E'$ where $v_1 \in V_1$ and $v_2 \in V_2$. Assume $L_{v_1} = \{a_1, \ldots, a_p\}$, $L_{v_2} = \{b_1, \ldots, b_q\}$. Note that for every $e \in E_2$ we have

$$\left|\{u = (e, b) \in \mathcal{U} \mid \exists v \in A_1 \cap \mathcal{S}', u \in v\}\right| \geq |B| - \frac{l|V'|^2}{|E|},$$

and for all $u = ((v_1, v_2), b) \in \mathcal{U}$, there exists $v \in A_1 \cap \mathcal{S}'$ where $u \in v$ iff $u \in S_{v_1, a_i}$ or $u \in S_{v_2, b_i}$. Thus we have

$$\left|(\{e\} \times B) \cap \left((\bigcup_{i=1}^{p} S_{v_1, a_i}) \cup (\bigcup_{j=1}^{q} S_{v_2, b_j})\right)\right| \geq |B| - \frac{l|V'|^2}{|E|},$$

which means

$$\left|(\bigcup_{i=1}^{p} C_{\Pi_e(a_i)}) \cup (\bigcup_{j=1}^{q} \overline{C_{b_j}})\right| \geq |B| - \frac{l|V'|^2}{|E|} = (1 - \frac{l|V'|^2}{|E||B|})|B| = (1 - (\frac{5}{147n})^r l)|B| > (1 - 2^{-l-1})|B|.$$

Thus by the definition of $(m, l, 2^{-l-1}))$-set system, we know that there exists i, j such that $\Pi_e(a_i) = b_j$. Therefore, e is satisfied with probability $\frac{1}{|L_{v_1}| \cdot |L_{v_2}|} \geq \frac{4}{l^2}$ because the labels are uniformly sampled. Thus the expected number of the edges satisfied in E' is at least $\frac{4}{l^2} \cdot \frac{|E|}{4} = \frac{|E|}{l^2}$, which means, there is a way to label all the vertices in V' and satisfies at least $\frac{|E|}{l^2}$ edges.

Finally, because $r = 3 \log \log n$ and $l = r \log n$, we know that at most $\varepsilon^r \cdot |E| < \frac{|E|}{l^2}$ edges can be satisfied by any labelling, which is a contradiction. □

With these lemmas, we can prove our lower bound on the TZ_2-optimization problem.

Proof of Theorem 2.5: By Lemma A.8 and Lemma A.9, we have a polynomial reduction from $\text{LabelCover}_{n,r,\varepsilon}$ problem to TZ_2-optimization problem, which maps a YES instance of $\text{LabelCover}_{n,r,\varepsilon}$ to a TZ_2-optimization instance with optimal cost at most $(|V'| + 1) \cdot |V|$, and maps a NO instance of $\text{LabelCover}_{n,r,\varepsilon}$ to a TZ_2-optimization instance with optimal cost at least $\frac{l}{8}|V'| \cdot |V|$. The gap is $\frac{\frac{l}{8}|V'| \cdot |V|}{(|V'|+1) \cdot |V|} = \Theta(\frac{l}{8}) = \Theta(\log |V|)$.

Combined with the hardness Theorem A.4, we know that unless $\mathbf{NP} \subseteq \mathbf{DTIME}(n^{O(\log \log n)})$, the TZ_2-optimization problem does not admit a polynomial-time $o(\log n)$-approximation. □

B Proofs in section 4 (integrality gap of TZ_k-optimization problem)

B.1 Relaxation validity

We first prove that our LP relaxation is indeed valid, i.e., we prove the following claim.

Claim B.1. *LP_{TZ_k} is a valid relaxation to the TZ_k-optimization problem.*

Proof. Let $A_1, \ldots, A_{k-1}$ be a valid solution to the TZ_k-optimization problem. Let $x_v^{(i)} = \mathbb{1}_{v \in A_i}$ and $y_{uv}^{(i)} = \mathbb{1}_{v \in R_{iu}}$ for all $i \in [k]$ and $u, v \in V$. We can see that the objective value $\sum_{i=1}^{k} \sum_{u,v \in V} y_{uv}^{(i)} = \sum_{i=1}^{k} \sum_{u \in V} |R_{iu}| = cost(A_1, \ldots, A_{k-1}, V, d)$, which is the cost function.

We can also see that the first constraint is satisfied by $x_v^{(i)}$ and $y_{uv}^{(i)}$ because $\varnothing = A_k \subseteq A_{k-1} \subseteq \ldots \subseteq A_0 = V$. The second constraint is satisfied because if $v \in A_{i-1}$ and there is no vertex in $A_i \cap B_u(v)$, then $v \in R_{iu}$. The third constraint is trivially satisfied.

Therefore $x_v^{(i)}$ and $y_{uv}^{(i)}$ is a valid solution to LP_{TZ_k} which makes the LP objective value equal to the actuall cost function. Thus the claim is proved. □

B.2 Integrality gap

Let's consider an instance (V, d) with $V = [n]$. All n vertices lie on a circle and they evenly split the cycle. The cycle distance $d(u, v) = \min\{|u - v|, n + \min\{u, v\} - \max\{u, v\}\}$.

We first show that on this instance, LP_{TZ_k} has a solution with low cost.

Lemma B.2. *LP_{TZ_k} has a solution with cost $O(n^{1+\frac{1}{2^{k-1}}})$ on instance (V, d).*

Proof. Consider the following setting of the LP variables: let $x_v^{(i)} = n^{-\frac{2^i-1}{2^{k-1}}}$ for all $v \in V$ and $i \in [k-1]$, and let $y_{uv}^{(i)} = \max\{0, x_v^{(i-1)} - \sum_{w \in B_u(v)} x_w^{(i)}\}$ for all $u, v \in V$ and $i \in [k]$.

We can see that $x_v^{(i)} = n^{-\frac{2^i-1}{2^{k-1}}} \geq n^{-\frac{2^{i+1}-1}{2^{k-1}}} = x_v^{(i+1)}$ which satisfies the first constraint of LP_{TZ_k}, $y_{uv}^{(i)} \geq x_v^{(i-1)} - \sum_{w \in B_u(v)} x_w^{(i)}$ which satisfies the second constraint of LP_{TZ_k}, and $y_{uv}^{(i)} \geq 0$ which satisfies the third constraint of LP_{TZ_k}. Therefore $x_v^{(i)}, y_{uv}^{(i)}$ is a valid solution to LP_{TZ_k}.

The objective value of this solution is

$$\sum_{i=1}^{k}\sum_{u,v\in V} y_{uv}^{(i)} = \sum_{i=1}^{k}\sum_{u,v\in V} \max\{0, x_v^{(i-1)} - \sum_{w\in B_u(v)} x_w^{(i)}\} \tag{1}$$

$$= \sum_{i=1}^{k-1}\sum_{u,v\in V} \max\{0, x_v^{(i-1)} - |B_u(v)| x_v^{(i)}\} + \sum_{u,v\in V}(x_v^{(k-1)} - 0) \tag{2}$$

$$= \sum_{i=1}^{k-1}\sum_{u,v\in V} \max\{0, n^{-\frac{2^{i-1}-1}{2^{k-1}}} - (2\cdot d(u,v)+1)\cdot n^{-\frac{2^i-1}{2^{k-1}}}\} + \sum_{u,v\in V} n^{-\frac{2^{k-1}-1}{2^{k-1}}} \tag{3}$$

$$= \sum_{i=1}^{k-1}\sum_{u\in V} \left(n^{-\frac{2^{i-1}-1}{2^{k-1}}} - n^{-\frac{2^i-1}{2^{k-1}}} + n^{-\frac{2^{i-1}-1}{2^{k-1}}} - 3\cdot n^{-\frac{2^i-1}{2^{k-1}}} + \dots\right) + n^{1+\frac{1}{2^{k-1}}} \tag{4}$$

$$= \sum_{i=1}^{k-1} n\cdot \left(O(n^{-\frac{2^{i-1}-1}{2^{k-1}}})\cdot O(n^{\frac{2^i-1}{2^{k-1}}-\frac{2^{i-1}-1}{2^{k-1}}})\right) + n^{1+\frac{1}{2^{k-1}}} \tag{5}$$

$$= \sum_{i=1}^{k-1} O(n^{1+\frac{1}{2^{k-1}}}) + n^{1+\frac{1}{2^{k-1}}} \tag{6}$$

$$= O(n^{1+\frac{1}{2^{k-1}}}) \tag{7}$$

Here equation (2) holds because of all $x_v^{(i)}$ are equal and all $x_v^{(k)} = 0$. Equation (3) holds because of the definition of circle distance. Equation (4) is a unrolling, and equation (5) is a summation over arithmetic progression. The last equation holds because of k is a constant. □

Next we will show that the optimal solution of this instance is large.

Lemma B.3. *The optimal solution to the instance (V, d) has cost at least $\Omega(n^{1+\frac{1}{k}})$.*

We will prove this lemma using a stronger claim. The lemma holds by setting $a = 1, b = \lfloor \frac{n}{2} \rfloor$, and $l = k$ in this claim:

Claim B.4. *For a segment $[a, b]$ of the cycle where $a, b \in [n]$, $b - a < \frac{n}{2}$, and all the vertices in $[a, b]$ are NOT in A_l, we have $\sum_{i=1}^{l}\sum_{u\in[a,b]\cap[n]} |R_{iu}| \geq \left(\frac{b-a+1}{4^l}\right)^{1+\frac{1}{l}}$ for each $l \in [k]$.*

Proof. We prove this by doing induction on l. The base case is $l = 1$. For each vertex $u \in [a, b]$, We know that

$$R_{1u} = \{v \in V \mid d(u,v) < \min_{w\in A_1} d(u,w)\} \subseteq \{v \in [a,b] \mid |u - v| \leq \min\{u - a, b - u\}\},$$

so $|R_{1u}| \geq 2\cdot\min\{u - a, b - u\}$ because all the vertices in $[a, b]$ are NOT in A_1. So

$$\sum_{u\in[a,b]} |R_{1u}| \geq 2\cdot(1 + 2 + \dots + \left\lfloor \frac{b-a+2}{2} \right\rfloor + \dots + 2 + 1) \geq \left(\frac{b-a+1}{4}\right)^2.$$

Now we consider general case $l \geq 2$, and assume the claim is established on $l - 1$.

Assume there are m vertices $t_1, \dots, t_m \in [a, \lceil \frac{a+b}{2} \rceil] \cap A_{l-1}$, and $[a, \lceil \frac{a+b}{2} \rceil]$ are splitted to small segments $[a_0, b_0], \dots, [a_m, b_m]$ where all the vertices in $[a_i, b_i]$ are not in A_{l-1} (if a segment has no

vertex inside, we let $b_i = a_i - 1$ without lose of generality). Then for each $i \in [m]$ and $u \in [a_i, b_i+1]$, we have $t_1, \ldots, t_i \in R_{lu}$ because $R_{lu} = \{v \in A_{l-1} \mid d(u,v) < \min_{w \in A_l} d(u,w)\}$. Thus

$$\begin{aligned}\sum_{i=1}^{l} \sum_{u \in [a,b] \cap [n]} |R_{iu}| &\geq \sum_{i=0}^{m} \left(\sum_{j=1}^{l-1} \sum_{u \in [a_i, b_i] \cap [n]} |R_{ju}| + \sum_{u \in [a_i, b_i] \cap [n]} |R_{lu}| \right) \\ &\geq \sum_{i=0}^{m} \left(\left(\frac{b_i - a_i + 1}{4^{l-1}} \right)^{1+\frac{1}{l-1}} + i \cdot (b_i - a_i + 2) \right).\end{aligned}$$

If $m > \frac{b-a+1}{4}$, $\sum_{i=0}^{m} i \cdot 1$ is already at least $\left(\frac{b-a+1}{4^l}\right)^{1+\frac{1}{l}}$.

If $m \leq \frac{b-a+1}{4}$, we have a stronger inequality which we will prove later:

Lemma B.5. *If $\alpha \in [1,2]$ and $x_i \geq 0$ for all $i \in [m]$, then*

$$\sum_{i=0}^{m} (x_i^{\alpha} + 4i \cdot x_i) \geq \left(\sum_{i=0}^{m} x_i \right)^{2-\frac{1}{\alpha}}$$

Using this inequality, by setting $x_i = \frac{b_i - a_i + 1}{4^{l-1}}$ and $\alpha = 1 + \frac{1}{l-1}$ we have

$$\begin{aligned}\sum_{i=1}^{l} \sum_{u \in [a,b] \cap [n]} |R_{iu}| &\geq \left(\sum_{i=0}^{m} \frac{b_i - a_i + 1}{4^{l-1}} \right)^{2-\frac{1}{1+\frac{1}{l-1}}} \\ &\geq \left(\frac{\lceil \frac{a+b}{2} \rceil - a + 1 - m}{4^{l-1}} \right)^{1+\frac{1}{l}} \\ &\geq \left(\frac{\frac{b-a}{2} + 1 - \frac{b-a}{4}}{4^{l-1}} \right)^{1+\frac{1}{l}} \geq \left(\frac{b-a+1}{4^l} \right)^{1+\frac{1}{l}}.\end{aligned}$$

□

With these lemma in hand, we can now prove Theorem 2.6.

Proof of Theorem 2.6:

Combine Lemma B.2 and Lemma B.3, there is an $\Omega(\frac{n^{1+\frac{1}{k}}}{n^{1+\frac{1}{2^k-1}}}) = \Omega(n^{\frac{1}{k} - \frac{1}{2^k-1}})$ integrality gap for the basic LP relaxation LP_{TZ_k}. □

Proof of Lemma B.5:

Let $M = (\sum_{i=1}^{m} x_i)^{\frac{\alpha-1}{\alpha}}$. We first split the problem to 2 cases, depending on whether $m \leq M$.

Case 1: $m \leq M$.

In this case, by Hölder's inequality, we have

$$(\sum_{i=0}^{m} x_i^{\alpha})^{\frac{1}{\alpha}} \cdot (\sum_{i=0}^{m} 1^{\frac{\alpha}{\alpha-1}})^{\frac{\alpha-1}{\alpha}} \geq (\sum_{i=0}^{m} x_i \cdot 1)$$

thus

$$\sum_{i=0}^{m} x_i^{\alpha} \geq \frac{(\sum_{i=0}^{m} x_i)^{\alpha}}{m^{\alpha-1}} \geq \frac{(\sum_{i=0}^{m} x_i)^{\alpha}}{M^{\alpha-1}} \geq (\sum_{i=0}^{m} x_i)^{\alpha - \frac{\alpha-1}{\alpha} \cdot (\alpha-1)} = (\sum_{i=0}^{m} x_i)^{2-\frac{1}{\alpha}}$$

Case 2: $m > M$.

Let's fix $T = \sum_{i=0}^{m} x_i$ and consider the $\mathbf{x}^*$ which minimizes the left side: $l(\mathbf{x}) = \sum_{i=0}^{m}(x_i^\alpha + 4i \cdot x_i)$.

Consider any two consecutive variables x_j and x_{j+1}, we claim that, in $\mathbf{x}^*$, for each $0 \le j < m$, either $x_{j+1}^* = 0$, or $(x_j^*)^{\alpha-1} - (x_{j+1}^*)^{\alpha-1} = \frac{4}{\alpha}$.

This is because, if we replace the x_{j+1} in $l(\mathbf{x})$ by $T - \sum_{i\neq(j+1)} x_i$ and do partial derivative with respect of x_j, we have

$$
\begin{aligned}
&\frac{\partial}{\partial x_j}\left(\sum_{i\neq(j+1)}(x_i^\alpha + 4i\cdot x_i) + \left(T - \sum_{i\neq(j+1)} x_i\right)^\alpha + 4(j+1)\cdot\left(T - \sum_{i\neq(j+1)} x_i\right)\right)\\
=&\frac{\partial}{\partial x_j}\left(x_j^\alpha + 4j\times x_j + \left(T - \sum\nolimits_{i\neq(j+1)} x_i\right)^\alpha + 4(j+1)\cdot\left(T - \sum\nolimits_{i\neq(j+1)} x_i\right)\right)\\
=&\alpha\cdot x_j^{\alpha-1} + 4j - \alpha\cdot\left(T - \sum\nolimits_{i\neq(j+1)} x_i\right)^{\alpha-1} - 4(j+1)\\
=&\alpha\cdot\left(x_j^{\alpha-1} - \left(T - \sum\nolimits_{i\neq(j+1)} x_i\right)^{\alpha-1}\right) - 4\\
=&\alpha(x_j^{\alpha-1} - x_{j+1}^{\alpha-1}) - 4.
\end{aligned}
$$

If we fix x_i for all $i \in [0,m] \cap \mathbb{N}\backslash\{j, j+1\}$, this partial derivative monotonically increases as x_j increases. Thus when $l(\mathbf{x})$ is minimized, either the partial derivative equals 0, which means $(x_j^*)^{\alpha-1} - (x_{j+1}^*)^{\alpha-1} = \frac{4}{\alpha}$, or x_j hits the ceiling, which means $x_j^* = T - \sum_{i\neq j,(j+1)} x_i^*$, so $x_{j+1}^* = 0$.

This result shows that, the number series $(x_0^*)^{\alpha-1}, (x_1^*)^{\alpha-1}, \ldots, (x_m^*)^{\alpha-1}$ is in decreasing order, where

$$
(x_i^*)^{\alpha-1} = \begin{cases} (x_{i-1}^*)^{\alpha-1} - \frac{4}{\alpha}, & \text{if } (x_i^*)^{\alpha-1} > \frac{4}{\alpha} \\ 0, & \text{otherwise} \end{cases}
$$

If the number of non-zero entries in $\mathbf{x}^*$ is at most M, then this comes back to the Case 1. Otherwise, there are more than M non-zero entries in $\mathbf{x}^*$, thus $x_0, x_1, \ldots, x_M$ are all non-zero, and $(x_i^*)^{\alpha-1} \ge \frac{4}{\alpha}\cdot(M-i)$ for $i \le M$. Therefore

$$
\begin{aligned}
\sum_{i=0}^{m}(x_i^*)^\alpha &\ge \sum_{i=0}^{M}\left(\frac{4}{\alpha}\cdot(M-i)\right)^{\frac{\alpha}{\alpha-1}} && (8)\\
&\ge \sum_{i=1}^{M}\left(\frac{4i}{\alpha}\right)^{\frac{\alpha}{\alpha-1}} && (9)\\
&\ge \frac{(\sum_{i=1}^{M}\frac{4i}{\alpha})^{\frac{\alpha}{\alpha-1}}}{M^{\frac{1}{\alpha-1}}} && (10)\\
&\ge \frac{(\frac{2M^2}{\alpha})^{\frac{\alpha}{\alpha-1}}}{M^{\frac{1}{\alpha-1}}} && (11)\\
&\ge (\frac{2}{\alpha})^{\frac{\alpha}{\alpha-1}}\cdot(\sum_{i=0}^{m}x_i)^{\frac{\alpha-1}{\alpha}\cdot(\frac{2\alpha}{\alpha-1}-\frac{1}{\alpha-1})} && (12)\\
&\ge (\sum_{i=0}^{m}x_i)^{2-\frac{1}{\alpha}} && (13)
\end{aligned}
$$

Here inequality (13) holds because of $\alpha \in [1,2]$. Inequality (10) holds because of Hölder's inequality

$$(\sum_{i=1}^{m} 1^{\alpha})^{\frac{1}{\alpha}} \cdot (\sum_{i=1}^{m} y_i^{\frac{\alpha}{\alpha-1}})^{\frac{\alpha-1}{\alpha}} \geq \sum_{i=1}^{m} y_i \cdot 1$$

□

C Proofs in section 5

C.1 Proof of Valid Relaxation

We prove the following claim:

Claim C.1. *LP_{PR} is a valid relaxation to the PR-optimization problem.*

Let A be a valid solution to the PR-optimization problem. Let $x_v = \mathbb{1}_{v \in A}$ and $y_{uv} = \mathbb{1}_{\{u,v\} \in R}$ for all $u, v \in V$. We can see that the objective value $\sum_{v \in V} n \cdot x_v + \sum_{\{u,v\} \subseteq V} y_{uv} = n \cdot |A| + |R| = cost(A, V, d)$, which is the cost function.

We can also see that the first constraint is satisfied by x_v and y_{uv} because if $y_{uv} = 0$, we have $d(u,v) \geq \min_{w \in A} d(u,w) + \min_{w \in A} d(v,w) - 1$, then for all $r \in [0, d(u,v)]$, there must be a vertex in $A \cap (B_u(r) \cup B_v(d(u,v) - r))$, which makes $0 \geq 1 - \sum_{w \in B_u(r) \cup B_v(d(u,v)-r)} x_w$ satisfied. The second and the third constraints are trivially satisfied.

Therefore x_v and y_{uv} is a valid solution to LP_{PR} which makes the LP objective value equal to the actuall cost function. Thus the claim is proved. □

C.2 Lower Bound Proofs

We start from the following theorem:

Theorem C.2 ([20])**.** *Unless $\mathbf{P} = \mathbf{NP}$, there is no $o(\log n)$-approximation to the set cover problem.*

We now prove two lemmas about the reduction (completeness and soundness).

Lemma C.3. *If there is a solution $\mathcal{S}^*$ to the set cover instance $(\mathcal{U}, \mathcal{S})$ where $|\mathcal{S}^*| = t$, then there is a set A where $cost(A, V, d) \leq t|V|$.*

Proof. For each $S \in S^*$, we add an arbitrary element from G_S to A. Then for every vertex in V, the closest vertex in A has distance at most 1 to it. Therefore

$$\begin{aligned} R &= \left\{ \{u,v\} \subseteq V \mid d(u,v) < \min_{w \in A} d(u,w) + \min_{w \in A} d(v,w) - 1 \right\} \\ &= \{\{u,v\} \subseteq V \mid d(u,v) < 1 + 1 - 1\} = \varnothing \end{aligned}$$

Thus the total cost is at most $|V| \cdot |A| + |R| = t|V|$. □

Lemma C.4. *If there is a set $A \subseteq V$ where $cost(A, V, d) \leq t|V|$, then there exists a solution $\mathcal{S}^*$ to the set cover instance $(\mathcal{U}, \mathcal{S})$ where $|\mathcal{S}^*| = t$.*

Proof. First, we say that a group $G = G_e$ or $G = G_S$ is "covered" if there exists a vertex $u \in G$, which $\min_{w \in A} d(u,w) = 1$. Then by the definition of d, it's easy to see that if a group G is covered, then for all vertices $u \in G$, we have $\min_{w \in A} d(u,w) = 1$. In addition, if a group G_e is covered, then either $G_e \cap A \neq \varnothing$, or there is a $S \in \mathcal{S}$, where $e \in S$ and $G_S \cap A \neq \varnothing$.

We can also see that, if a group G is not covered, then let

$$\begin{aligned} R_G =& \left\{ \{u,v\} \subseteq G \mid d(u,v) < \min_{w\in A} d(u,w) + \min_{w\in A} d(v,w) - 1 \right\} \\ =& \{\{u,v\} \subseteq G \mid d(u,v) < 2+2-1\} \\ =& \{\{u,v\} \subseteq G\} \end{aligned}$$

Thus $|R_G| \geq \frac{3n(3n-1)}{2} > 3n^2 > |V|$. Therefore if we add an arbitrary element from G to A, then $|R|$ decreases by at least $|R_G| \geq |V|$, and $|V| \cdot |A|$ increases by $|V|$, which makes $cost(A,V,d)$ only decrease. If we keep doing this operation, there will be a set A which makes sure that all the groups are covered, and $cost(A,V,d) \leq t|V|$. Now, for every vertex $u \in V$, we have $\min_{w\in A} d(u,w) = 1$, so $R = \varnothing$.

We can keep modifying A to the form we want. If there is a vertex $v \in G_e \cap A$, removing v and simultaneously adding a vertex in any $S \ni e$ to A does not increase the cost. This is because this operation keeps the fact that all the groups are covered.

Finally, we have a set A where only contains vertices in $\bigcup_{S\in\mathcal{S}} G_S$ and $cost(A,V,d) \leq t|V|$. Let $\mathcal{S}^* = \{S \in \mathcal{S} \mid G_S \cap A \neq \varnothing\}$. Then $|\mathcal{S}^*| \leq t$ because $cost(A,V,d) = |A| \cdot |V| + |R| \geq |\mathcal{S}^*| \cdot |V|$, and $\mathcal{S}^*$ covers $\mathcal{U}$ because all the group G_e are covered. □

These lemmas, combined with Theorem C.2, imply Theorem 2.8

D Proofs in section 6

D.1 TZ_2-Optimization Problem With Outliers

D.1.1 Proof of expected cost

Lemma D.1. *If $\|\vec{y}^*_{uv}\|^2 \leq \frac{\varepsilon}{2}$, then the probability that $uv \in R_{1u}$ is at most $\frac{1}{n}$.*

Proof. If $\|\vec{z}^*_u\|^2 \geq \frac{1}{1+\varepsilon}$ or $\|\vec{z}^*_v\|^2 \geq \frac{1}{1+\varepsilon}$, then u or v is in F, so $v \notin R_{1u}$. Thus we only consider the case that $\|\vec{z}^*_u\|^2 < \frac{1}{1+\varepsilon}$ and $\|\vec{z}^*_v\|^2 < \frac{1}{1+\varepsilon}$, which means $\vec{z}^*_u \cdot \vec{z}^*_v < \frac{1}{1+\varepsilon}$. Since $\|\vec{y}^*_{uv}\|^2 \leq \frac{\varepsilon}{2}$, we have

$$\sum_{w\in B_u(v)} \|\vec{x}^*_w\|^2 \geq 1 - \frac{\varepsilon}{2} - \frac{1}{1+\varepsilon} \geq \frac{\varepsilon}{3}.$$

Therefore, the probability that $d(u,v) < \min_{w\in A} d(u,w)$ is at most

$$\prod_{w\in B_u(v)} (1 - \min\{\frac{3\ln n}{\varepsilon} \cdot \|\vec{x}^*_v\|^2, 1\}) \leq e^{-\sum_{w\in B_u(v)} \frac{3\ln n}{\varepsilon} \cdot \|\vec{x}^*_v\|^2} \leq \frac{1}{n}$$

□

Therefore, let $OPT_{SDP_{TZ_2O}}$ denotes the optimal cost of SDP_{TZ_2O}, then the expected cost of the rounding algorithm is at most

$$\sum_{v\in V} (n-f) \cdot \frac{3\ln n}{\varepsilon} \cdot \|\vec{x}^*_v\|^2 + \frac{2}{\varepsilon} \cdot \sum_{u,v\in V} \|\vec{y}^*_{uv}\|^2 + n^2 \cdot \frac{1}{n} \leq O(\log n) \cdot SDP_{TZ_2O} + n \leq O(\log n) \cdot OPT$$

because $OPT \geq \Omega(n)$, which proves Theorem 2.9.

D.1.2 True approximation

When the number of outliers is low, in particular when $f \le \sqrt{n}$, we can find an actual $O(\log n)$-approximation.

The SDP and rounding algorithm are the same, except we will choose f vertices with the highest $\|\vec{z}_v^*\|^2$ values as F, rather than a threshold rounding of $\frac{1}{1+\varepsilon}$.

Now there are two cases when $\|\vec{y}_{uv}^*\|^2 \le \frac{\varepsilon}{2}$. One case is the same as before, where $\sum_{w\in B_u(v)} \|\vec{x}_w^*\|^2 \ge \frac{\varepsilon}{3}$. In this case, the probability that $v \in R_{1u}$ is at most $\frac{1}{n}$. The other case is $\sum_{w\in B_u(v)} \|\vec{x}_w^*\|^2 < \frac{\varepsilon}{3}$, which means $\vec{z}_u^* \cdot \vec{z}_v^* \ge 1 - \frac{\varepsilon}{2} - \frac{\varepsilon}{3} = 1 - \frac{5}{6}\varepsilon$.

However, this case will not appear a lot. Whenever $\vec{z}_u^* \cdot \vec{z}_v^* \ge 1 - \frac{5}{6}\varepsilon$, both $\|\vec{z}_u^*\|$ and $\|\vec{z}_v^*\|$ should be at least $1 - \frac{5}{6}\varepsilon$, which means $\|\vec{z}_u^*\|^2$ and $\|\vec{z}_v^*\|^2$ is at least $\frac{1}{2}$. Because $\sum_{v\in V} \|\vec{z}_v^*\|^2 \le f$, we know that there are at most $2f$ of $\|\vec{z}_v^*\|^2$ are at least $\frac{1}{2}$. Therefore the number of u, v pairs that $\|\vec{y}_{uv}^*\|^2 \le \frac{\varepsilon}{2}$ and $\sum_{w\in B_u(v)} \|\vec{x}_w^*\|^2 < \frac{\varepsilon}{3}$ is at most $2f \cdot 2f = 4n$.

Therefore, let $OPT_{SDP_{TZ_2O}}$ denotes the optimal cost of SDP_{TZ_2O}, then the expected cost of the rounding algorithm is at most

$$\sum_{v\in V}(n-f)\cdot\frac{3\ln n}{\varepsilon}\cdot\|\vec{x}_v^*\|^2+\frac{2}{\varepsilon}\cdot\sum_{u,v\in V}\|\vec{y}_{uv}^*\|^2+n^2\cdot\frac{1}{n}+4n \le O(\log n)\cdot OPT_{SDP_{TZ_2O}}+5n \le O(\log n)\cdot OPT$$

because $OPT \ge \Omega(n)$, which proves Theorem 2.10.

D.2 PR-Optimization Problem With Outliers

Lemma D.2. *If $\|\vec{y}_{uv}^*\|^2 \le \frac{\varepsilon}{2}$, then the probability that $\{u, v\} \in R$ is at most $\frac{1}{n}$.*

Proof. If $\|\vec{z}_u^*\|^2 \ge \frac{1}{1+\varepsilon}$ or $\|\vec{z}_v^*\|^2 \ge \frac{1}{1+\varepsilon}$, then u or v is in F, so $\{u, v\} \notin R$. Thus we only consider the case that $\|\vec{z}_u^*\|^2 < \frac{1}{1+\varepsilon}$ and $\|\vec{z}_v^*\|^2 < \frac{1}{1+\varepsilon}$, which means $\vec{z}_u^* \cdot \vec{z}_v^* < \frac{1}{1+\varepsilon}$. Since $\|\vec{y}_{uv}^*\|^2 \le \frac{\varepsilon}{2}$, we have

$$\sum_{w\in B_u(r)\cup B_v(d(u,v)-r)} \|\vec{x}_w^*\|^2 \ge 1 - \frac{\varepsilon}{2} - \frac{1}{1+\varepsilon} \ge \frac{\varepsilon}{3}.$$

Therefore, the probability that $A \cap (B_u(r) \cup B_v(d(u, v) - r)) = \varnothing$ for a specifiic $r \in [0, d(u, v)]$ is at most

$$\prod_{w\in B_u(r)\cup B_v(d(u,v)-r)} (1 - \min\{\frac{6\ln n}{\varepsilon} \cdot \|\vec{x}_w^*\|^2, 1\}) \le e^{-\sum_{w\in B_u(v)} \frac{6\ln n}{\varepsilon}\cdot\|\vec{x}_w^*\|^2} \le \frac{1}{n^2}$$

□

By using union bound over all the different r we used in our SDP, the probability that there exists an $r \in [0, d(u, v)]$ where $A \cap (B_u(r) \cup B_v(d(u, v) - r)) = \varnothing$ is at most $\frac{1}{n^2} \cdot n = \frac{1}{n}$, which means $d(u, v) < \min_{w\in A} d(u, w) + \min_{w\in A} d(v, w) - 1$ with probability at most $\frac{1}{n}$, so the probability that $\{u, v\} \in R$ is at most $\frac{1}{n}$.

Therefore, let $OPT_{SDP_{PR}}$ denotes the optimal cost of SDP_{PR}, then the expected cost of the rounding algorithm is at most

$$\sum_{v\in V}(n-f)\cdot\frac{3\ln n}{\varepsilon}\cdot\|\vec{x}_v^*\|^2+\frac{2}{\varepsilon}\cdot\sum_{u,v\in V}\|\vec{y}_{uv}^*\|^2+n^2\cdot\frac{1}{n} \le O(\log n)\cdot OPT_{SDP_{PR}}+n \le O(\log n)\cdot OPT$$

because $OPT \ge \Omega(n)$, which proves Theorem 2.11.